arXiv:1702.06280v1 [cs.CR] 21 Feb 2017

On the (Statistical) Detection of Adversarial Examples

Kathrin Grosse[†]*, Praveen Manoharan*[†]*, Nicolas Papernot*[‡]*, Michael Backes*[†]*, Patrick McDaniel*[‡]
CISPA, Saarland University[†]*; Penn State University*[‡]
grosse@cs.uni-saarland.de

Abstract

Machine Learning (ML) models are applied in a variety of tasks such as network intrusion detection or malware classification. Yet, these models are vulnerable to a class of malicious inputs known as adversarial examples. These are slightly perturbed inputs that are classified incorrectly by the ML model. The mitigation of these adversarial inputs remains an open problem.

As a step towards a model-agnostic defense against adversarial examples, we show that they are not drawn from the same distribution than the original data, and can thus be detected using statistical tests. As the number of malicious points included in samples presented to the test diminishes, its detection confidence decreases. Hence, we introduce a complimentary approach to identify specific inputs that are adversarial among sets of inputs flagged by the statistical test. Specifically, we augment our ML model with an additional output, in which the model is trained to classify all adversarial inputs.

We evaluate our approach on multiple adversarial example crafting methods (including the fast gradient sign and Jacobian-based saliency map methods) with several datasets. The statistical test flags sample sets containing adversarial inputs with confidence above 80%. Furthermore, our augmented model either detects adversarial examples with high accuracy ($> 80\%$) or increases the adversary's cost—the perturbation added—by more than 150%. In this way, we show that statistical properties of adversarial examples are essential to their detection.

1 Introduction

Machine learning algorithms are usually designed under the assumption that models are trained on samples drawn from a distribution that is representative of test samples for which they will later make predictions—ideally, the training and test distributions should be identical. However, this does not hold in the presence of adversaries. A motivated adversary may either manipulate the training [3] or test [29] distribution of a ML system. This has severe consequences when ML is applied to security-critical problems. Attacks are increasingly elaborate, as demonstrated by the variety of strategies available to evade malware detection built with ML [29, 33].

Often, adversaries construct their attack inputs from a benign ML input. For instance, the feature vector of a malware—correctly classified by a ML model as malware—can be modified into a new feature vector, the *adversarial example*, that is classified as benign [12, 33].

Defenses proposed to mitigate adversarial examples, such as adversarial training [9] and defensive distillation [25], all fail to adapt to changes in the attack strategy. They both make it harder for the adversary to craft adversarial examples using existing techniques only, thus creating an arms race [6, 21]. However, we argue that this arms race is not inevitable: by definition, adversarial examples must exhibit some statistical differences with the legitimate data on which ML models perform well.

Hence, we develop in this work a countermeasure that uses the distinguishability of adversarial examples with respect to the expected data distribution. We use statistical testing to evaluate the hypothesis that adversarial examples, crafted to evade a ML model, are outside of the training distribution. We show that the approach is well-suited and performs well on diverse datasets [17, 2, 1] and adversarial example algorithms [9, 23, 22].

However, this test needs to be presented with a sufficiently large sample set of suspicious inputs—as its confidence diminishes with the number of malicious inputs in the sample set. Therefore, we propose a second complimentary mechanism for detecting individual adversarial examples. The idea also exploits the statistical distinguishability of adversarial examples to design an outlier detection system, but this time it is directly integrated in the ML model. Indeed, we show that models can be augmented with an additional output reserved for adversarial examples—in essence training the model with adversar-

ial examples as their own class. The model, trained to map all adversarial inputs to the added output, exhibits robustness to adversaries.

Our contributions are the following:

Statistical Test—In Section 5, we design a statistical test to distinguish adversarial examples from the model's training data. Among tests proposed in the literature, we select the kernel-based two-sample test introduced by Gretton et al. [11]. This test has the key benefit of being model-agnostic; because its kernel allows us to apply the test directly on samples from the ML model's input data.

We demonstrate the good performance of this test on three datasets: MNIST (hand-written digits), DREBIN (Android malware) and MicroRNA (medical data). Specifically, we show that the test can confidently detect samples of 50 adversarial inputs when they differ from the expected dataset distribution. Results are consistent across multiple generation techniques for adversarial examples, including the fast gradient sign method [9] and the Jacobian-based saliency map approach [23]. Furthermore, the confidence of the statistical test remains high even when analyzing a mixture of adversarial examples generated with multiple algorithms.

Integrated Outlier Detection—As the statistical test's confidence diminishes when it is presented with increasingly small sample sets of adversarial inputs, we propose another outlier detection system. We add an additional class to the model's output, and train the model to recognize adversarial examples as part of this new class. The intuition behind the idea is the same (detect adversarial examples using their statistical properties) but this approach allows the defender to detect individual adversarial examples among a set of inputs identified as malicious (by the statistical test for instance).

We observe here that adversarial examples lie in unexpected regions of the model's output surface because they are not representative of the distribution. By training the model to identify out-of-distribution inputs, one removes at least part of its error away by filling in the things that are demonstrable (e.g., adversarial examples).

In Section 6, we find that this approach correctly assigns adversarial examples as being part of the outlier class with over 80% success for two of the three datasets considered. For the third dataset, they are not frequently detected but the perturbation that an adversary needs to add to mislead the model is increased by 150%. Thus, the cost of conducting an attack is increased in all cases.

Arms race—We then investigate adversarial strategies taking into account the defense deployed. For instance, black-box attacks were previously shown to evade adversarial training and defensive distillation [21]. The adversary uses an (undefended) auxiliary model to find adversarial inputs that are also misclassified by the defended model (because the defended model makes adversarial crafting harder but does not solve the model error). Our mechanisms perform well under such black-box scenarios: adversarial inputs crafted by a black-box attack are more likely to be detected than those computed directly by an adversary with access to our model.

2 Background

We provide here the relevant background on ML and adversarial ML. We finally give an overview of the statistical hypothesis test applied in this paper.

2.1 Machine Learning Classifiers

We introduce ML notation used throughout this paper. All ML models considered are classifiers and learn a function $f(x) \mapsto y$. An input point or example $x \in X$ is made up of n components or features (e.g., all system level calls made by an Android application), and $y \in Y$ is a label (e.g., malware or benign). In classification problems, the possible values of y are discrete. The output of the model, however, is often real valued probabilities over the set of possible labels, from which the most likely label is inferred as the one with the largest probability.

In other words, there is an underlying and almost always unknown distribution $D^{C_i}_{\mathsf{real}}$ for each class C_i. The set of training data X is sampled from this distribution, and the classifier approximates this distribution during training, thereby learning $D^{C_i}_{\mathsf{train}}$. The set of test data X_t, used to validate the classifier's performance, is assumed to be drawn from the same $D^{C_i}_{\mathsf{real}}$.

Next, we present typical ML models used to solve classification problems and studied in this paper.

Decision Trees—These models are composed of internal nodes and leafs, whose graph makes up a tree. Each leaf is assigned one of the possible labels, while the intermediate nodes form a path of conditions defined using the input features. An example is classified by finding a path of appropriate conditions from the root to one of the leaves. Decision trees are created by successively maximizing the information gain resulting from the choice of a condition as a way to partition the data in two subsets (according to the value of an input feature).

Support Vector Machines—They compute a $n-1$ dimensional hyperplane to separate the training points. Since there are infinitely many such hyperplanes, the one with the largest margin is computed—yielding a convex optimization problem given the training data (X,Y).

Neural Networks—They are composed of small computational units called *neurons* that apply an *activation function* to their weighted input. Neurons are organized in interconnected layers. Depending on the number of

layers, a network is said to be *shallow* (single intermediate layer) or *deep* (several intermediate layers). Information is propagated through the network by having the output of a given layer be the input of the following layer. Each of these links is parameterized by weights. The set of model weights—or model parameters—are trained to minimize the model's prediction error $\|f(x)-y\|$ on a collection of known input-output pairs (X,Y).

Logistic regression—This linear model can be conceptualized as a special case of neural networks without hidden layers. For problems with two classes, the logistic function is the activation function. For multi-class problems, it is the softmax. They are trained like neural nets.

2.2 Adversarial Machine Learning

Adversarial ML [15], and more generally the security and privacy of ML [24], is concerned with the study of vulnerabilities that arise when ML is deployed in the presence of malicious individuals. Different attack vectors are available to adversaries. They can target ML during training [3] of the model parameters or during test time [29] when making predictions.

In this paper, we defend against test time attacks. They target a trained model $f(_,\theta)$, and typically aim to find an example x' similar to an original example x, which is however classified differently. To achieve this, a perturbation δ with same dimensionality as x is computed:

$$f(x',\theta) \neq f(x,\theta) \text{ where } x' = x+\delta \text{ and } \min \delta$$

where δ is chosen to be minimal to prevent detection and as to indirectly represent the attackers limitations when perturbing features. When targeting computer vision, the perturbation must not be detectable to the human eye. When targeting a malware detector, the perturbation must not remove the application's malicious behavior. Instead of simply having inputs classified in a wrong class, the attacker can also target a particular class.

A typical example of such attacks is the evasion of a bayesian spam filter, first demonstrated by Lowd et al. [20]. Malware detection systems have also been targeted, as shown by Srndic et al. [29] or Grosse et al. [12]. In addition, these adversarial inputs are known to transfer across (i.e., to mislead) multiple models simultaneously [30]. This transferability property was used to create attacks against black-box ML systems in settings where the adversary has no access to the model or training data [21, 22]. A detailed discussion of some of these attacks can be found in Section 4.

Several defenses for attacks at test time have been proposed. For instance, training on adversarial inputs pro-actively [9] or performing defensive distillation [25]. Both of them may fail due to gradient masking [21]. Other approaches make use of game theory [7, 19, 5]. However, they are computationally expensive.

2.3 Statistical Hypothesis Testing

The framework of two-sample statistical hypothesis testing was introduced to determine whether two randomly drawn samples originate from the same distribution.

Formally, let $X \sim p$ denote that sample X was drawn from a distribution p. A statistical test can then be formalized as follows: let $X_1 \sim p$, where $|X_1| = n$ and $X_2 \sim q$, where $|X_2| = m$. The null hypothesis H_0 states that $p = q$. The alternative hypothesis, H_A, on the other hand, is that $p \neq q$. The statistical test $\mathscr{T}(X_1,X_2) : \mathscr{X}^n \times \mathscr{X}^m \to \{0,1\}$ takes both samples as its input and distinguishes between H_0 and H_A. In particular, the statistic returned is matched to a p-value, refering to the probability that we obtain the observed outcome or a more extreme one. The p-value is fixed before the application of the test, typically at 0.05 or 0.01. If the test statistic is smaller than the p-value, H_0 is rejected. A consistent test will reject H_0 when $p \neq q$ in the large sample size limit.

There are several two-sample tests for higher dimensions. For instance, the Hotellings T^2 test evaluates whether two distributions have the same mean [14]. Several other tests depending on graph or tree properties of the data were proposed by Friedman et al. [8], Rosenbaum et al. [26] or Hall et al. [13].

Most of these tests are not appropriate when considering data with high dimensionality. This led Gretton et al. [11] to introduce a kernel-based test. In this case, we measure the distance between two probabilities (represented by samples X_1 and X_2). In practice, this distance is formalized as the biased estimator of the true Maximum Mean Discrepancy (MMD):

$$\text{MMD}_b[\mathscr{F},X_1,X_2] = \sup_{f\in\mathscr{F}} \left(\frac{1}{n}\sum_{i=1}^{n} f(x_{1i}) - \frac{1}{m}\sum_{i=1}^{m} f(x_{2i}) \right)$$

where the maximum indicates that we pick the kernel function f from the function class $\mathscr{F}$ that maximize the difference between the functions. Further, in contrast to other measures, we do not need the explicit probabilities.

Gretton et al. [11] introduced several tests. We focus on one of them in the following: a test based on the asymptotic distribution of the unbiased MMD. However, to consistently estimate the distribution of the MMD under H_0, we need to bootstrap[1]. Here, bootstrapping refers to a subsampling method, where one samples from the data available with replacement. By repeating this procedure many times, we obtain an estimate for the MMD value under H_0.

[1]Other methods have been proposed, such as moment matching Pearson curves. We focus here on one specific test used in this paper.

3 Methodology

Here, we introduce a threat model to characterize the adversaries our system is facing. We also derive a formal argument justifying the statistical divergence of adversarial examples from benign training points. This observation underlies the design of our defensive mechanisms.

3.1 Threat model

Adversarial knowledge— Adversarial example crafting algorithms proposed in the literature primarily differ in the assumptions they make about the knowledge available to adversaries [24]. Algorithms fall in two classes of assumptions: *white-box* and *black-box*.

Adversaries operating in the white-box threat model have unfettered access to the ML system's architecture, the value of its parameters, and its training data. In contrast, other adversaries do not have access to this information. They operate in a black-box threat model where they typically can interact with the model only through an interface analog to a cryptographic oracle: it returns the label or probability vector output by the model when presented with an input chosen by the adversary.

In this paper, we are designing a defensive mechanism. As such, we must consider the worst-case scenario of the strongest adversary. We therefore operate in both the white-box and the black-box threat model. While our attacks may not be practical for certain ML systems, it allows us to provide stronger defensive guarantees.

Adversarial capabilities— These are only restricted by constraints on the perturbations introduced to craft adversarial examples from legitimate inputs. Such constraints vary from dataset to dataset, and as such we leave their discussion to the description of our setup in Section 4.

3.2 Statistical Properties of Adversarial Examples

When learning a classifier from training data as described in Section 2, one seeks to learn the real distributions of features $D^{C_i}_{\mathsf{real}}$ for each subset C_i corresponding to a class i. These subsets define a partition of the training data, i.e. $\cup_i C_i = X$. However, due to the limited number of training examples, any machine learning algorithm will only be able to learn an approximation of this real distribution, the *learned feature distributions* $D^{C_i}_{\mathsf{train}}$.

A notable result in ML is that any *stable* learning algorithms will learn the real distribution $D^{C_i}_{\mathsf{real}}$ up to any multiplicative factor given a sufficient number of training examples drawn from $D^{C_i}_{\mathsf{real}}$ [4]. Stability refers here to the fact that given a slight modification of the data, the resulting classifier and its prediction do not change much. Coming back to our previous reasoning, however, this *full generalization* is in practice impossible due to the finite (and often small) number of training examples available.

The existence of adversarial examples is a manifestation of the difference between the real feature distribution $D^{C_i}_{\mathsf{real}}$ and the learned feature distribution $D^{C_i}_{\mathsf{train}}$: the adversary follows the strategy of finding a sample drawn from $D^{C_i}_{\mathsf{real}}$ that does not adhere to the learned distribution $D^{C_i}_{\mathsf{train}}$. This is only partially dependent on the actual algorithm used to compute the adversarial example. Yet, the adversary (or any entity as a matter of fact) does not know the real feature distribution $D^{C_i}_{\mathsf{real}}$ (otherwise one could use that distribution in lieu of the ML model). Therefore, existing crafting algorithms generate adversarial examples by perturbing legitimate examples drawn from $D^{C_i}_{\mathsf{train}}$, as discussed in Section 2.2.

Independently of how adversarial examples were generated, all adversarial examples for a class C_i will constitute a new distribution $D^{C_i}_{\mathsf{adv}}$ of this class. Following the above arguments, clearly $D^{C_i}_{\mathsf{adv}}$ is consistent with $D^{C_i}_{\mathsf{real}}$, since each adversarial example for a class C_i is still a data point that belongs to this class. On the other hand, however, $D^{C_i}_{\mathsf{adv}} \neq D^{C_i}_{\mathsf{train}}$. This follows from a reductio ad absurdum: if the opposite was true, adversarial examples would be correctly classified by the classifier.

As discussed in Section 2.3, consistent statistical tests can be used to detect whether two sets or samples X_1 and X_2 were sampled from the same distribution or not. A sufficient (possibly infinite) number of examples in each sample allows such a consistent statistical test to detect the difference in the distributions even if the underlying distributions of X_1 and X_2 are very similar.

Following from the above, statistical tests are natural candidates for adversarial example detection. Adversarial examples have to inherently be distributed differently from legitimate examples used during training. The difference in distribution should consequently be detectable by a statistical test. Hence, the first hypothesis we want to validate or invalidate is the ability of a statistical test to distinguish between benign and adversarial data points. We have two practical limitations, one is that we can do so by observing a finite (and small) number of examples, the second that we are restricted to existing adversarial example crafting algorithms.

Hypothesis 1. *We only need a bounded number of n examples to observe a measurable difference in the distribution of examples drawn* D_{adv} *and* D_{train} *using a consistent statistical test* T.

We validate this hypothesis in Section 5. We show that as few as 50 misclassified adversarial examples per class are sufficient to observe a measurable difference between legitimate trainings points and adversarial examples for existing adversarial example crafting algorithms.

3.3 Detecting Adversarial Examples

The main limitation of statistical tests is that they cannot detect adversarial examples on a per-input basis. Thus, the defender must be able to collect a sufficiently large batch of adversarial inputs before it can detect the presence of adversaries. In addition, the defender cannot identify specific inputs that were manipulated by the adversary among batches of examples sampled. Indeed, sampling a single example will not allow us to confidently estimate its distribution with a statistical test.

This may not be acceptable in security-critical applications. A statistical test, itself, is therefore not always suitable as a defensive mechanism.

However, we propose another approach to leverage the fact that D_{adv} is different from D_{train}. We augment our learning model with an additional *outlier class* C_{out}. We then train the ML model to classify adversarial examples in that class. Technically, C_{out} thus contains all examples that are not drawn from any of the learned distributions $D^{C_i}_{\mathsf{train}}$. We seek to show that this augmented classifier can detect newly crafted adversarial examples at test time.

Hypothesis 2. *The augmented classifier with an outlier class C_{out} successfully detects adversarial examples.*

We validate Hypothesis 2 in Section 6. We also address the potential existence of an arms race between attackers and defenders in Section 7.

4 Experimental setup

We describe here the experimental setup used in Sections 5, 6 and 7 to validate the hypotheses stated in Section 3. Specifically, we design our setup to answer the following experimental questions:

- ***Q1: In practice, how well do statistical tests distinguish adversarial distributions from legitimate ones?*** In Section 5, we first find that the MMD and energy distance can statistical distinguish adversarial examples from legitimate inputs. Statistical tests can thus be designed based on these metrics to detect adversarial examples crafted with several known techniques. In fact, we find that often a sample size of 50 is enough to identify them.
- ***Q2: Can detection be integrated in ML models to identify individual adversarial examples?*** In Section 6, we show that a neural network trained with an additional outlier class detects $> 80\%$ of the adversarial examples it is presented with. We also find that in cases where malicious inputs are undetected, the perturbation introduced to evade the model needs to be increased by 150%, making the attack more expensive for attackers.
- ***Q3: Do our defenses create an arms race?*** Our statistical test performs well even when the sample set only contains a small subset of adversarial inputs. We also find that the model with an outlier class is robust to adaptive adversaries, such as the ones using black-box attacks. Even when such adversaries are capable of closely mimicking our model to perform a black-box attack, they are still detected with 60% accuracy in the worst case, and in many cases with accuracies larger than $> 90\%$.

To answer these questions comprehensively, we use datasets, models and adversarial example algorithms representative of the ML space. We now describe them.

4.1 Adversarial example crafting

In our experiments, we consider the following attacks.

Fast Gradient Sign Method (FGSM)—This attack computes the gradient of the model's output with respect to its input. It then perturbs examples in that direction. The computational efficiency of this attack comes at the expense of it introducing large perturbations that affect the entire input. We used the implementation provided in the `cleverhans` library [10].

Jacobian-based Saliency Map Approach (JSMA)—In contrast to the FGSM, this attack iteratively computes the best feature to perturb. This yields an adversarial example with fewer modified features, at the expense however of a higher computational cost. We rely again on the implementation provided in the `cleverhans` library [10].

SVM attack—This attack is described in [22]. It targets a linear SVM by shifting the point orthogonally along the decision boundary. The result is a perturbation similar to the one found by the FGSM. In the case of SVMs however, the perturbation depends on the target class.

Decision Tree (DT) attack—We implemented a variant of the attack from [22] where we search the shortest path between the leaf in which the sample is currently at and the closest leaf of another class. We then perturb the feature that is used in the first common node shared by the two paths. By repeating this process, we achieve a misclassification. This attack modifies only few features.

4.2 Datasets

We evaluate our hypothesis on three datasets.

MNIST—This dataset consists of black-and-white images from 0 to 9 taking real values [17]. It is composed of $60,000$ images, of which $10,000$ form a test dataset. Each image has $28x28$ pixels.

DREBIN—This malware dataset contains $545,333$ binary malware features [2]. To make adversarial example crafting faster, we apply dimensionality reduction, as

done by Grosse et al. [12], to obtain 955 features. The dataset contains 129,013 Android applications, of which 123,453 are benign and 5,560 are malicious. We split this dataset randomly in training and test data, where the test data contains one tenth of all samples.

Due to its binary nature, it is straightforward to detect attacks like the FGSM or the SVM attack: they lead to non-binary features. We did, nonetheless, compute them in several settings to investigate performance of the detection capabilities. We did not restrict the features that can be perturbed (in contrast to previous work[12]) in an effort to evaluate against stronger adversaries.

MicroRNA—This medical dataset consists of 3966 samples, of which 1280 are breast cancer serums and the remaining are non-cancer control serums. We restrict the features to the 5 features reported as most useful by the original authors [1]. When needed, we split the dataset randomly in training and test subsets, with a $1/10$ ratio. This dataset contains real-valued features, each with different mean and variance. We computed perturbations (for SVM, the FGSM and the JSMA) dependent on the variance of the feature to be perturbed.

4.3 Models

We now describe the details of the models used. These models were already introduced in Section 2.1.

Decision Trees—We use the Gini impurity as the information gain metric to evaluate the split criterion.

Support Vector Machines—We use a linear multi-class SVM. We train it with an l2 penalty and the squared hinge loss. When there are more then two classes, we follow the one-vs-rest strategy.

Neural Networks—For MNIST, the model has two convolutional layers, with filters of size $5x5$, each followed by max pooling. A fully connected layer with 1024 neurons follows. The DREBIN model reproduces the one described in [12]. It has two fully connected layers with 200 neurons each. The network on the MicroRNA data has a single hidden layer with 4 neurons.

All activation functions are ReLU. All models are further trained using early stopping and dropout, two common techniques to regularize the ML model's parameters and thus improve its generalization capabilities when the model makes predictions on test data.

Logistic regression—We train a logistic regression on MicroRNA data with dropout and a cross-entropy loss.

Most of these classifiers achieve accuracy comparable to the state-of-the art on MNIST[2]. For DREBIN, the accuracy is larger than 97.5%. On MicroRNA, the neural network and logistic regression achieve 95% accuracy.

[2]The linear SVM and decision tree only achieve 92.7% and 67.4%

5 Identifying adversarial examples using statistical metrics and tests

In the following, we conduct a series of experiments to answer the first question introduced in Section 4: *in practice, how well do statistical tests distinguish adversarial distributions from legitimate ones?* We find that two statistical metrics, the MMD and the energy distance, both reflect changes—that adversarial examples make to the underlying statistical properties of the distribution—by often strong variations of their value. Armed with these metrics, we apply a statistical test. It detects adversarial examples confidently, even when presented with small sample sets. This validates Hypothesis 1 from Section 3: adversarial examples exhibit statistical properties significantly different from legitimate data.

5.1 Characterizing adversarial examples with statistical metrics

We consider two statistical distance measures commonly used to compare higher dimensional data: (1) the maximum mean discrepancy, and (2) the energy distance.

Maximum Mean Discrepancy (MMD)— Recall from Section 2.3 that this divergence measure is defined as:

$$\mathrm{MMD}_b[\mathscr{F}, X_1, X_2] = \sup_{f \in \mathscr{F}} \left(\frac{1}{n} \sum_{i=1}^{n} f(x_{1i}) - \frac{1}{m} \sum_{i=1}^{m} f(x_{2i}) \right)$$

where $x_{1i} \in X_1$ is the i-th data point in the first sample. $x_{2j} \in X_2$ is the j-th data point in the second sample, which is possibly drawn from another distribution than X_1. $f \in \mathscr{F}$ is a kernel function chosen to maximize the distances between the samples from the two distributions. In our case, a Gaussian kernel is used.

Energy distance (ED)— The ED, which is also used to compare the statistical distance between two distributions, was first introduced by Szkély et al. [31]. It is a specific case of the maximum mean discrepancy, where one does not apply any kernel.

Measurement results— We perturb the training distribution of several MNIST models (a neural network, decision tree, and support vector machine) using the appropriate adversarial example crafting algorithms presented in Section 4.1. We then measure the statistical divergence (i.e. the distance) between the adversarially manipulated training data and the model's training data by computing the MMD and ED.

All data points are drawn randomly out of the 60,000 training points. The chance of having a particular sample and its modified counterpart in the same batch is thus very small. Measurements are reported in Table 1. Note that we also give the distance between the unmodified training and test distributions to provide a baseline.

Manipulation	Parameters	MMD	ED
Original	-	*130.85*	*0.105*
FGSM	$\varepsilon = 0.07$	157.904	0.281
FGSM	$\varepsilon = 0.275$	213.967	0.603
JSMA	-	137.63	0.14
DT attack	-	130.71	0.1
SVM attack	$\varepsilon = 0.25$	186.32	0.524
Flipped	-	135.0	0.306
Subsampling	45 pixel	102.7	2.159
Gaussian Blur	4 pixel	128.52	1.021

Table 1: MMD and ED between the original distribution and manipulated distributions obtained by several adversarial and geometric techniques on MNIST. Values are averaged over random samples of size 1,000. For each technique, parameters such as the perturbation magnitude for the FGSM or the number of blurred pixels are given. The JSMA leads to a change of on average 20 pixels, whereas the DT attack changes on average 1 pixel.

We observe that for most adversarial examples, there is a strong increase in values of the MMD and ED. In the case of the FGSM, we observe that the increase is stronger with larger perturbations ε. For the JSMA and the DT attack, changes are more subtle because these approaches only modify very few features.

We then manipulate the test data using geometric perturbations. While these are not adversarial, they are nevertheless helpful to interpret the magnitude of the statistical divergences. Perturbations considered consist in mirroring the sample, subsampling from the original values, and introducing Gaussian blur.[3] We find that mirroring and subsampling affect both the MMD and ED, whereas Gaussian blur only significantly increases the ED.

In this first experiment, we observed that there exists measurable statistical distances between benign and malicious inputs. Therefore, we investigate next whether one can design consistent statistical tests to confidently identify samples from adversarial distributions.

5.2 Detecting adversarial examples using hypothesis testing

We apply a statistical test to evaluate the following hypothesis: *samples from the test distribution are statistically close to samples from the training distribution*. We expect this hypothesis to be accepted for samples from the legitimate test distribution, but rejected for samples containing adversarial examples. Indeed, we observed in Section 5.1 that adversarial distributions statistically diverge from the training distribution.

[3]This geometric perturbation approximates an attack against ML models introduced by Biggio et al. [3]. Indeed, the adversarial inputs produced by this attack appear as blurry, with less crisp shapes.

Two-sample hypothesis testing— As stated before, the test we chose is appropriate to handle high dimensional inputs and small sample sizes.[4] We compute the biased estimate of MMD using a Gaussian kernel, and then apply 10,000 bootstrapping iterations to estimate the distributions. Based on this, we compute the test statistic and compare it to the p-value. We report all results given a p-value of 0.05. For samples of legitimate data, this statistic should always be very high, whereas for sample sets containing adversarial examples, we expect it to be low—since they are sampled from a different distribution and thus the hypothesis should be rejected.

The test is more likely to detect a difference in two distributions when it considers samples of large size (i.e., the sample contains more inputs from the distribution).

Whenever we write *confidently detected* at sample size x, we mean that all 200 instances of the test on x randomly sampled examples from each of the two distributions rejected H_0. Percentages reported correspond to the n times of 200 the test accepted H_0.

Results— Regardless of the sample size, the hypothesis acceptance for benign data generally lie around 95%. This means that the benign data is confidently identified as such. The sample size required to confidently detect adversarial distributions is given in Table 2 for the three datasets (MNIST, DREBIN, MicroRNA).

For most datasets and models, a sample size of 50 adversarial examples is sufficient for the statistical test to reject H_0 when comparing this sample to a sample from the benign distribution. Thus, the statistical test identifies adversarial examples with strong confidence, despite having few points of comparison (relatively to the training set size of 50,000 for MNIST).

Some exceptions should be noted. A sample of at least 100 inputs is required to confidently detect adversarial examples crafted with the JSMA on a neural network. In addition, the test is unable to detect adversarial examples crafted for MNIST on decision trees.

However, these two observations are consistent with results from Section 5.1, which showed that these attacks yielded adversarial inputs with less distinguishable statistical properties than the FGSM for instance. Another result confirms this explanation: the SVM attack, which was observed to lead to large changes in the MMD and energy distance, is as well easily detected by the two-sample test. A sample of 10 adversarial inputs is sufficient to confidently reject the H_0.

To summarize, we empirically showed that, for a representative collection of datasets, it is possible to tell apart JSMA and FSGM adversarial examples from legitimate data using statistical hypothesis tests.

[4]We used an implementation publicly available at https://github.com/emanuele/kernel_two_sample_test.

Dataset	FGSM	JSMA	SVM at.	DT at.
MNIST	50 (.275)	100 (16)	10 (.25)	- (1)
DREBIN	50 (.6)	50 (2)	10 (.25)	50 (2)
Micro	50 (.6)	10 (3)	*	50 (1)

Table 2: Smallest sample size needed to confidently detect adversarial examples. The average adversarial perturbation introduced is characterized either by stating in parenthesis the perturbation parameter ε (FGSM, SVM attack), or the number of perturbed features (JSMA, DT attack). Stars indicate experiments that were not conducted due to the technique's ineffectiveness (it yielded reductions in accuracy smaller than $< 30\%$).

Dataset	FGSM	JSMA	SVM at.	DT at.
MNIST	50 (50)	50 (100)	10 (10)	50 (-)
DREBIN	10 (10)	50 (50)	10 (10)	50 (-)
DREBIN	10 (10)	50 (50)	10 (10)	10 (50)
Micro	10 (10)	10 (10)	*	10 (50)
Micro	10 (10)	10 (10)	*	10 (10)

Table 3: Smallest sample size needed to confidently detect adversarial examples when using labels as additional information. Label information is either original label (in brackets) or class that is wrongly assigned by the classifier. Perturbations are equivalent to Table 2. For MNIST, we report the largest sample sizes needed for all classes. First entry for DREBIN refers to malware class, second to benign programs. For MicroRNA, upper row is the control group, second are the cancer serums.

5.3 Decreasing sample size by performing class-wise statistical tests

In an effort to reduce the number of inputs (i.e., the sample size) required for the statistical test to be confident, we propose here to refine it. We previously assumed that the dataset was generated by a single underlying distribution. However, each class of the problem can be viewed as its own distribution. We thus perform below the statistical test on samples grouped according to their label. We separate the inputs according to their original label, or the (wrong) label assigned by the model. These tests are found to be confident for smaller sample sizes.

Results—Experimental results are given in Table 3. On MNIST, we find that these new class-wise statistical tests reduce the sample size needed to detect the JSMA to 50 examples. This is is also the case with DREBIN, where the sample size is reduced to 10. For the MicroRNA we only observe a change in sample size concerning the decision tree attack, where again confident detection is already possible at a sample size of 10.

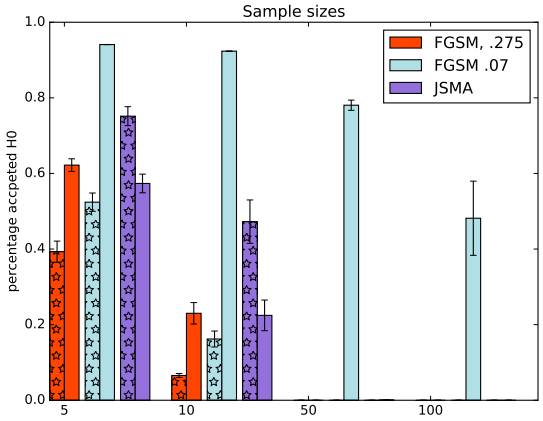

Figure 1: Class wise test on MNIST. Given a sample set size (x-axis), we report how frequently H_0 is accepted by the kernel test on average (y-axis, lower is better). Label information is either original label (patterns) or class that is wrongly assigned by the classifier (plain).

For all datasets, we find that using a statistical test based on the distribution of the class in which the inputs are (wrongly) misclassified is more effective than using the class from which they were derived: the latter even completely fails for decision trees on MNIST. These results are consistent across the two other datasets (DREBIN and MicroRNA). In Figure 1, however, we observe that on MNIST when testing for the FGSM examples, this tendency is reversed.

Briefly put, we observed that the minimum sample size to achieve confident detection with a class-wise test is smaller than the sample size required by the general statistical test. We also noted that statistical tests comparing adversarial examples with training examples from the class they are misclassified as (rather than the class they were derived from) are more confident.

6 Integrating Outlier Detection in Models

In the previous section, we concluded that the distribution of adversarial examples statistically differs from the expected distribution. Hence, our statistical test can detect the presence of an adversary. Yet, its confidence diminishes with the number of examples in the sample set analyzed: this test cannot be used to identify which specific inputs are adversarial among a set of inputs.

In this Section, we provide an answer to our second experimental question: "Can we detect individual adversarial examples?" Our approach adds an additional output to the model. The model is trained to assign this new output class to all adversarial inputs. In other words, we

explicitly train models to label all inputs that are not part of the expected distribution as part of a new *outlier class*.

In the following experiments, we show that this approach is complimentary to the statistical test introduced in Section 2.3 because it enables the defender to accurately identify whether a given input is adversarial or not.

Training with an outlier class— We start the process by training an initial model N_{in} on the original data $D = \{X, Y\}$. We compute adversarial examples for N_{in} on the training data, denoted as X_{in}. We then train a new model, N_{p1} on an augmented dataset, $X \cup X_{in}$, where all adversarial examples are assigned to the outlier class. Specifically, we arrange batches of inputs analyzed by the learning algorithm such that $2/3$ are legitimate training examples from X and $1/3$ are adversarial from X_{in}.

To evaluate the robustness of this model, we craft adversarial examples on the test data $D_t = \{X_t, Y_t\}$ for the augmented model. We assume an adversarial example to be correctly classified if it is classified in its original class or as an outlier, and report this accuracy. Additionally, we report the percentage of all adversarial examples that were classified as outliers.

Performance impact— In order for the outlier class to be a practical mechanism for detecting adversarial inputs, the impact on the model accuracy on benign test data should be limited.

On MNIST, we observe that the parameters of our model augmented with an outlier class converge slightly slower (in terms of training iterations required) to their final values compared to an unmodified model. In addition, the model's accuracy varies between 98.0% and 98.8% when it is presented with legitimate inputs. This decrease from the unmodified model's 99.2% accuracy is the cost of robustness to adversaries.

On DREBIN, the original network achieves an accuracy on the benign test data of 97.7%. Considering the mean, we observe a slight decrease of 0.2% in accuracy for the model trained using an additional class.

On the microRNA data, the original logistic regression achieves an accuracy around 95.4%. With an outlier class, accuracy drops to 93.4%. The same holds for neural networks trained on the same dataset.

Detection results— In this series of experiments, we first measure the detection performance of the model on malicious examples. We then study the failures cases on benign data of our mechanism.

The results on MNIST are depicted in Table 4. The model detects 99.92% of the FGSM ($\varepsilon = 0.275$) adversarial examples and 83.78% of the JSMA ($\varepsilon < 80$). As a consequence of our model being trained with an outlier class, we also observed that the mean number of components perturbed by the saliency map approach increased from 16.1 to 22.2. We also observe that the model generalizes to various parameters of the attack, e.g. when trained using FGSM examples setting $\varepsilon = 0.275$, the network is as well robust to attacks using other values of ε. In particular, the network is more robust if the ε is higher.

ε attack	accuracy	detected outliers
0.2	99.8%	99.78%
0.275	99.92%	99.91%
0.3	99.94%	99.93%
0.4	99.99%	99.99%
0.5, 0.6	100%	100%
≤ 80	83.78%	83.76%

Table 4: Accuracy on MNIST of recomputed adversarial examples. Examples for training using $\varepsilon = 0.275$ for the FGSM and $\varepsilon \leq 200$ for the JSMA.

network	accuracy	outliers	mod. feat.
common	0.01%	0%	2.3
JSMA	0.01%	0%	5.8

Table 5: Accuracy on DREBIN of recomputed adversarial examples. Parameters were $\varepsilon \leq 90$ for the JSMA both for training and the recomputed attack.

The results on DREBIN are summarized in Table 5. Concerning the FGSM adversarial examples, we observed, independently from chosen ε, an accuracy around 7.7%. The network could further not be hardened against those adversarial examples: after training, the accuracy was still 7.8%. For the JSMA, we observe an initial accuracy of 0.009% by changing 2.3 features.[5] When retraining on adversarial examples, we do not observe any increase in robustness. We do observe, however, an increase in the number of changed features up to 5.8 when trained on JSMA examples.

To understand whether the limited effectiveness of our defense on DREBIN is a consequence of the binary nature of its data or the stronger success of the attack, we implemented a second attack with a worse heuristic that initially modifies 6.6 features on average. By training on adversarial examples crafted with this defense, it became impossible to craft adversarial examples using the same modified JSMA with an upper limit of 90 changed features.

We also trained a simple logistic regression on the MicroRNA data and trained it on adversarial examples. The results are depicted in Table 6. We only applied the FGSM, since the JSMA was not successful (only 40% of the adversarial examples evaded the model). Initially, we could decrease the accuracy to 4.3% by applying a

[5]This initial perturbation is much less then reported in the original work, since we do not restrict the features as done previously [12].

	common	JSMA	log reg	JSMA	NN
ε	acc.	acc.	out	acc.	out
0.2	87.4%	96.5%	16.7%	95.5%	21.9%
0.4	42.9%	94.1%	64.6%	89.6%	69.9%
0.6	15.9%	68.7%	64.9%	95.5%	92.6%
0.8	5.1%	30.8%	26.3%	98%	96%
1.0	4.3%	12.8%	6.7%	99.3%	97.5%

Table 6: Accuracy of a logistic regression and neural network on MicroRNA when presented with recomputed adversarial examples. Training was performed with the FGSM at $\varepsilon = 1.0$ and $\varepsilon = 0.8$. First column indicates the FGSM parameter value used to recompute adversarial examples after training is completed.

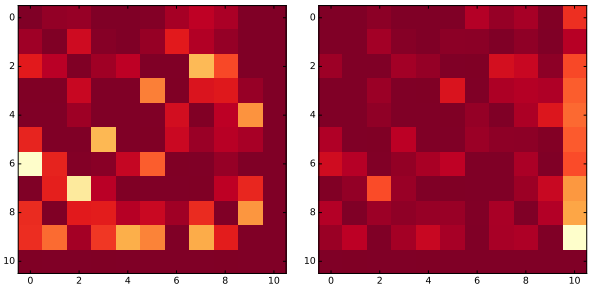

Figure 2: Confusion matrices on benign test sets of MNIST. The horizontal axis denotes the original label, and the vertical one the output class of the network. Left side corresponds to training on FGSM examples, right side on JSMA examples. The diagonal (correctly classified data points) has been zeroed, indices correspond to MNIST classes, where 10 is the outlier class. Both matrices are normalized in the same scale, i.e., same color means same number of misclassified samples. Brighter indicates a higher number of misclassified samples.

perturbation of $\varepsilon = 1.0$. Further increase of ε had no effect. Training logistic regression with an outlier class on adversarial examples given $\varepsilon = 1.0$, we achieve an accuracy which is 12.8%. For lower perturbations, we can increase the accuracy, however. The limited improvement is most likely due to the limited capacity of logistic regression models, which prevent them from learning models robust to adversarial examples [9]. Thus, we trained a neural network on the data. Initially it could be evaded with the same perturbation magnitude and success. Yet, when trained with the outlier class, we achieved an accuracy of 99.3% on the strongest ε.

Wrongly classified benign test data— Next, we investigate the error cases of our mechanism. First, the number of false positives, benign test examples of the original data that are wrongly classified as outliers, represents a small percentage of inputs: e.g., 0.5% on MNIST. In addition, we draw confusion matrices for the benign test data in Figure 2. The diagonal indicates correctly classified examples and is canceled out to better visualize out-of-diagonal and misclassified inputs. Misclassification between classes is very similar in the original case and when training with the FGSM. We thus omit the confusion matrix for original data in the interest of space. In contrast, when training on JSMA examples, a large fraction of misclassified data points is no longer misclassified as a legitimate class, but wrongly classified as outliers.

7 Preventing the Arms Race

A key challenge in ML security lies in the fact that no defense guarantees resilience to future attack designs. This contrasts with ML privacy where differential privacy guarantees withstand all hypothetical adversaries. Such an arms race may only be broken by mechanisms that have been proven to be secure in an expressive security model, such as the one of differential privacy.

While providing any formal guarantees for the methods proposed here is intrinsically hard given the nature of optimization problems solved by ML algorithms, we evaluate here their resilience to adaptive strategies. We first show that the statistical test still performs well when presented with a mixture of benign and malicious inputs. We then demonstrate the robustness of our models augmented with the outlier class to powerful black-box strategies that have evaded previous defenses.

7.1 Robustness of the Statistical Test

Assume we perform the statistical test before submitting inputs to the ML model. A natural question is whether an adversary aware of this defense could evade it by constructing adversarial examples simultaneously misleading the model and the statistical test.

Theoretically, a statistical test cannot be misled because adversarial examples necessarily deviate from the expected distribution (see Section 3). Yet, this assumes that the defender is capable of running statistical tests on sufficiently large sample sets of inputs. In that case, it is guaranteed that the null hypothesis would be rejected. However, this may not always be the case in practice as such sampling may require a potentially infinite number of inputs. As such, we measure the confidence of our statistical test as it is presented with more realistic sample sets of inputs in an effort to demonstrate its robustness.

We consider two scenarios were the attacker adapted its strategy. First, the attacker might hide small numbers of adversarial examples among a large number of benign samples. In the following, we thus investigate the statistical test's performance in detecting adversarial inputs

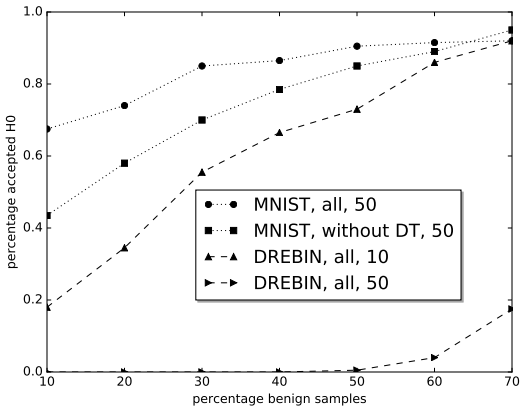

Figure 3: Evaluating mixtures of different adversarial examples and benign data. X-axis indicates the percentage of benign data. Y-axis is the percentage where the statistical test confirms that the data is from the same distribution (lower is better). We observe that the more benign examples (moving to the right), the harder it is to detect the remaining adversarial examples.

in the presence of legitimate inputs. Second, we consider sample sets of adversarial examples, where the adversary executed more than one adversarial crafting algorithm, like for instance the FGSM and the JSMA.

Mixture of adversarial examples—We observe reductions in the detection confidence for mixtures of adversarial examples. This is especially true when one of the adversarial example type used is itself hard to detect. For instance, we observed in Section 5.2 that adversarial inputs for decision trees are hard to detect on MNIST. This reduces the performance of the statistical test on samples that contain these examples. Our full results are depicted in Figure 3. Unless stated otherwise, the same mixtures of adversarial examples were used than previously.

Mixtures of adversarial and legitimate inputs—The trade-off between the ratio of benign and adversarial examples, test confidence and sample size is shown in Figure 3. The test is more confident when the percentage of adversarial examples is high or the sample size is large. Hence, one is less likely to detect adversarial examples mixed with legitimate inputs among small sets of inputs.

We conclude that the statistical test's confidence decreases, as expected, when the adversary submits very few adversarial examples among large sets of legitimate inputs. This is however the main motivation behind the outlier class mechanism introduced in Section 6.

7.2 Robustness of the Outlier Class

We now investigate the performance of models augmented with an outlier class in the face of adversaries aware of that defense. We first show that these models are able to generalize to varying attacker strategies, i.e., they can detect adversarial inputs crafted using a different algorithm than the one used to train the outlier class.

In addition, we not that in all previous experiments, we considered adversaries directly computing adversarial examples based on the defended model's parameters. Instead, we here evaluate the model's robustness when attacked using black-box strategies. These powerful attacks have been shown to evade previously proposed defenses, such as adversarial training and defensive distillation [21]. The reason is that these defenses did not actually fix model errors but rather manipulated the model's gradients, thus making it harder for the adversary to craft adversarial examples directly on the model.

A simple but highly successful strategy is then for the adversary to train an auxiliary model that mimics the defended model's predictions, and then use the auxiliary model to find adversarial examples that are also misclassified by the defended model. In the following, we show that our models with an outlier class are also robust to such strategies.

Robustness of detection to adaptive attackers— We investigate whether the outlier class generalizes to other adversarial example crafting techniques. In other words, we ask whether defending against one type of adversarial examples is sufficient to mitigate an *adaptive* attacker using multiple techniques to craft adversarial examples.

We thus trained the MNIST model's outlier class with only one kind of adversarial example, and then observed its robustness to another kind of adversarial examples.Table 7 reports this result for varying attack parameter intensities. If we train the model on JSMA adversarial examples, it is also robust to adversarial FGSM examples crafted. If perturbations are high, $> 97\%$ percent are correctly classified. The reverse case, however, does not hold: a model trained on FGSM is only slightly more robust then the original model.

We did not perform this experiment on DREBIN or MicroRNA datasets, since we could only apply one attack on each of them.

Robustness of detection to black-box attacks performed using transferability— We now show that our proposed outlier class mechanism is robust to an additional attack vector against ML models: black-box attacks exploiting adversarial example transferability. These techniques allow an adversary to force a ML model to misclassify without knowledge of its model parameters (and sometimes even without knowledge of its training data) by computing adversarial examples on a

Training on	ε attack	accuracy	detected outliers
$\varepsilon \leq 200$	0.1	79.2%	77.16%
$\varepsilon \leq 200$	0.275	97.05%	96.6%
$\varepsilon \leq 200$	0.4	98.67%	98.45%
$\varepsilon \leq 200$	0.6	99.71%	99.58%
$\varepsilon = 0.275$	≤ 80	9.63%	9.63%

Table 7: Evaluation of the outlier class on MNIST. The class is trained on adversarial examples crafted with either $\varepsilon = 0.275$ for the FGSM or $\varepsilon \leq 200$ for the JSMA. The detected outliers indicate how frequently the adversarial examples computed on the approach *not* used during training are detected at test time.

different model than the one targeted [30, 9, 21].

In order to simulate the *worst-case* adversary, we train the *substitute* model from which we will transfer adversarial examples back to the victim model with the same training data than the one used by the victim. Again, a black-box attack may be possible even without such knowledge [21]. Yet, this allows us to consider a particularly strong adversary capable to closely mimic our model. We further focus on the robustness of our approach to such attacks, rather than a demonstration of the feasibility of such attacks. We train two substitute models, one including the outlier class (denoted BB+1) as does the victim, the other not (BB). We then compute adversarial examples on the substitute model and evaluate the accuracy of the victim model's prediction.

MNIST results— The results on MNIST are given in Table 8. We observe in both settings high robustness against the adversarial examples computed using the FSGM. Depending on the setting, we even detect 100% of those examples when $\varepsilon > 0.4$ for BB+1 and $\varepsilon > 0.275$ for BB. We further observe high robustness as well concerning the adversarial examples computed using JSMA: In the setting of BB+1, the accuracy is around 93.3%, for BB even 97.3%. In all cases, most of the adversarial examples are classified as outliers.

DREBIN results— Results on DREBIN are in Table 9. We observe that, though the network is vulnerable to direct attacks, it is much more resistant to adversarial examples crafted on another network: We observe an accuracy around 60% on adversarial examples crafted using the JSMA. For JSMA examples crafted on the original network, accuracy is even higher and around 99%.

MicroRNA results— The results on the MicroRNA data are depicted in Table 10. We observe that logistic regression works well in detecting smaller perturbations till 0.4. For larger ε, however, the accuracy drastically decreases, partially to 12%. Again, we assign the small improvement to the limited capacity of logistic regression, we trained additionally neural networks. We observed higher accuracy in settings where the perturbation is maximmal ($\varepsilon = 1.0$). In general, we observe that logistic regression is more robust in the BB setting, whereas neural networks are more robust in the BB+1 setting (both with some exceptions).

	BB		BB+1	
ε attack	accuracy	outliers	accuracy	outliers
0.1	98.39%	55.72%	99.28%	81.64%
0.275	100%	100%	99.98%	99.98%
> 4	100%	100%	100%	100%
≤ 80	97.3%	97%	93.31%	92.85%

Table 8: Robustness of the outlier class to black-box attacks on MNIST. The outlier detection rates are computed by crafting adversarial examples on a substitute model, which is either trained wihtout (BB) or with an additional class (BB+1). The outlier class of the model targeted by the attack was trained on adversarial examples generated with the FGSM and JSMA.

substitute	ε attack	accuracy	outliers
BB	$\varepsilon < 90$	99.2%	98.6%
BB+1	$\varepsilon < 90$	64.2%	47.9%

Table 9: Robustness of the outlier class on DREBIN. Adversarial examples are crafted using the JSMA a substitute model which is either trained without (BB) or with an additional class (BB+1). The outlier class of the substitute and targeted model are trained using the JSMA.

8 Discussion

We discuss the limitations of the mechanisms proposed to detect adversarial examples: statistical testing, and an outlier class. We also explore avenues for future work.

Statistical Test— As we have seen, one of the major strengths of kernel-based statistical tests it that they operate and thus detect the presence of adversarial examples already in feature space, before these inputs are even fed to the ML model. Intuitively, we observed that the larger the perturbation applied is, the more likely it is to be confidently detected by the statistical test. Adversarial example crafting techniques that modify few features (like the JSMA or the decision tree attack) or perturb the features only slightly (small values of ε for the FGSM) are less likely to be detected.

This finding is consistent with the underlying stationary assumption made by all ML approaches. Since adversarial examples are not drawn from the same distribu-

ε	BB (reg)	BB (NN)	BB+1 (reg)	BB+1 (NN)
0.2	96.5 (11.9)	96.5 (12.9)	96.2 (14.9)	96.5 (15.7)
0.4	96.4 (46.7)	96.4 (48.7)	94.9 (59.8)	95 (65.2)
0.6	96.2 (74.7)	96 (80.6)	83.5 (74.5)	73.4 (68.9)
0.8	90.2 (75.3)	89.1 (81.8)	61.8 (55.8)	34.8 (30.8)
1.0	77.5 (68.2)	69.2 (65.4)	42.9 (38.6)	12 (6.0)
0.2	97.2 (13.9)	97.2 (14.6)	97.8 (21.5)	96.2 (12.1)
0.4	98.2 (54.0)	98.2 (55.3)	98.7 (68.4)	96 (45.7)
0.6	89.4 (76.0)	95 (78.3)	97.2 (93.2)	96.2 (81.6)
0.8	89.4 (76.0)	86.3 (80.3)	96.2 (94.2)	98.2 (95.0)
1.0	82.8 (79.8)	75.2 (73.2)	96.7 (94.7)	99.2 (97.2)

Table 10: Robustness of the outlier class on MicroRNA models. The adversarial examples are crafted using the FGSM. Both the target and substitute models are logistic regression (upper) or neural networks (below). The substitute model is either trained without (BB) or with an additional class (BB+1).

tion than benign data, the classifier is incapable of classifying them correctly. This property also holds for the training data itself, and as such, we expect our technique to generalize to the detection of poisoning attacks. In such attacks, the adversary attempts to degrade learning by inserting malicious points in the model's training data. This is however outside the scope of this work, and we leave this question to future work.

Integrating Outlier Detection— We further observe that adding an outlier class to the model yields robustness to adaptive attack strategies. For instance, on MNIST, training on adversarial examples computed using the JSMA hardens against computing FGSM examples. Perhaps surprisingly, this does not hold the other way around, which may suggest that the FGSM moves inputs further away from legitimate classes—in a different half space, whereas the JSMA would move adversarial examples closer to other classes. The confusion matrices from Figure 2 seem to support that hypothesis: the outlier class trained on JSMA examples contains benign data points whereas the FGSM one does not.

Future work should also investigate whether the robustness of the outlier class depends on the data itself: for instance, we suspect it may be influence by inter-class and intra-class distances. In particular, we observed that for DREBIN the outlier class was not as effective. This may well be a consequence of the binary nature of the data or close proximity of the two classes (malware and benign), as indicated by the very small perturbation magnitudes needed to mislead a model.

9 Related Work

Other approaches to detect malicious data points by statistical means have been proposed. However, they all depend on some of the internal activations of deep neural networks models [18, 28]. Hence, these approaches only apply to the specific classifier studied. In contrast, we apply our statistical test directly in feature space, allowing us to propose a model-agnostic defense.

Wang et al. [32] present a similar formal intuition as we do (using an oracle instead of the underlying distribution, though). From this, they formally derive conditions when a classifier is secured against adversarial examples. Further, they proposed a modified version of adversarial training, originally introduced by Goodfellow et al. [9]. In contrast to both of these approaches, we classify adversarial examples in a separate (and additional) outlier class. We also do not compute adversarial examples throughout training but rather use adversarial examples precomputed on a different model before training.

Metzen et al. [16] augment neural networks with an auxiliary network used to detect malicious samples. This additional network shares some of its parameters with the original one. In contrast, our outlier class mechanism is applicable to any ML models. In addition, their approach is limited in settings where the adversary adapts its strategy. Instead, our experiments systematically explore the space of adversaries (with different adversarial example crafting algorithms, datasets and models). We also present a detailed discussion of possible adaptive strategies, such as powerful black-box attacks known to be hard to defend against [21].

10 Conclusion

We empirically validated the hypothesis that adversarial examples can be detected using statistical tests before they are even fed to the ML model as inputs. Thus, their malicious properties are model-agnostic. In future work, we plan to investigate how these findings can be applied to design anomaly detection systems for ML inputs.

Furthermore, we show how to augment ML models with an additional class in which the model is trained to classify all adversarial inputs. This results in robustness to adversaries, even those using attack strategies based on transferability—a class of attacks known to be harder to defend against than gradient-based strategies.

Additionally, we expect that combining our approaches together, as well as with other defenses may prove beneficial. For instance, we expect defensive distillation and the statistical test or outlier class to work well together, as defensive distillation has been found to increase the perturbations that an adversary introduces.

Acknowledgments

Nicolas Papernot is supported by a Google PhD Fellowship in Security. Research was supported in part by the Army Research Laboratory, under Cooperative Agreement Number W911NF-13-2-0045 (ARL Cyber Security CRA), and the Army Research Office under grant W911NF-13-1-0421. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Army Research Laboratory or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for government purposes notwithstanding any copyright notation hereon.

References

[1] A. SHIMOMURA, SHIINO S, K. J. T. S. E. A. Novel combination of serum microrna for detecting breast cancer in the early stage. *Cancer Sci Mar 107(3):326-34* (2016).

[2] ARP, D., SPREITZENBARTH, M., HUBNER, M., GASCON, H., AND RIECK, K. DREBIN: Effective and Explainable Detection of Android Malware in Your Pocket. In *Proceedings of the 2014 Network and Distributed System Security Symposium (NDSS)* (2014).

[3] BIGGIO, B., NELSON, B., AND LASKOV, P. Poisoning attacks against support vector machines. *arXiv preprint arXiv:1206.6389* (2012).

[4] BOUSQUET, O., AND ELISSEEFF, A. Stability and generalization. *The Journal of Machine Learning Research 2* (2002), 499–526.

[5] BRÜCKNER, M., AND SCHEFFER, T. Stackelberg games for adversarial prediction problems. In *Proceedings of the 17th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, San Diego, CA, USA, August 21-24, 2011* (2011), pp. 547–555.

[6] CARLINI, N., AND WAGNER, D. Towards evaluating the robustness of neural networks. *CoRR abs/1608.04644* (2016).

[7] DALVI, N., DOMINGOS, P., SANGHAI, S., VERMA, D., ET AL. Adversarial classification. In *Proceedings of the tenth ACM SIGKDD international conference on Knowledge discovery and data mining* (2004), ACM, pp. 99–108.

[8] FRIEDMAN, J., AND RAFSKY, L. Multivariate generalizations of the wald–wolfowitz and smirnov two-sample tests. *Ann. Stat.; (United States) 7:4* (Jan 1979).

[9] GOODFELLOW, I. J., ET AL. Explaining and harnessing adversarial examples. In *Proceedings of the 2015 International Conference on Learning Representations* (2015).

[10] GOODFELLOW, I. J., PAPERNOT, N., AND MCDANIEL, P. D. cleverhans v0.1: an adversarial machine learning library. *CoRR abs/1610.00768* (2016).

[11] GRETTON, A., BORGWARDT, K. M., RASCH, M. J., SCHÖLKOPF, B., AND SMOLA, A. J. A kernel two-sample test. *Journal of Machine Learning Research 13* (2012), 723–773.

[12] GROSSE, K., PAPERNOT, N., MANOHARAN, P., BACKES, M., AND MCDANIEL, P. Adversarial perturbations against deep neural networks for malware classification. *CoRR abs/1606.04435* (2016).

[13] HALL, P., AND TAJVIDI, N. Permutation tests for equality of distributions in high-dimensional settings. *Biometrika 89*, 2 (2002), 359–374.

[14] HOTELLING, H. The generalization of student's ratio. *Ann. Math. Statist. 2*, 3 (08 1931), 360–378.

[15] HUANG, L., JOSEPH, A. D., NELSON, B., RUBINSTEIN, B. I. P., AND TYGAR, J. D. Adversarial machine learning. In *Proceedings of the 4th ACM Workshop on Security and Artificial Intelligence, AISec 2011, Chicago, IL, USA, October 21, 2011* (2011), pp. 43–58.

[16] JAN HENDRICK METZEN, TIM GENEWEIN, V. F., AND BISCHOFF, B. On detecting adversarial perturbations. (to appear).

[17] LECUN, Y., BOTTOU, L., BENGIO, Y., AND HAFFNER, P. Gradient-based learning applied to document recognition. *Proceedings of the IEEE 86*, 11 (November 1998), 2278–2324.

[18] LI, X., AND LI, F. Adversarial examples detection in deep networks with convolutional filter statistics. *CoRR abs/1612.07767* (2016).

[19] LIU, W., AND CHAWLA, S. Mining adversarial patterns via regularized loss minimization. *Machine Learning 81*, 1 (2010), 69–83.

[20] LOWD, D., AND MEEK, C. Good word attacks on statistical spam filters. In *CEAS 2005 - Second*

Conference on Email and Anti-Spam, July 21-22, 2005, Stanford University, California, USA (2005).

[21] PAPERNOT, N., MCDANIEL, P., GOODFELLOW, I., JHA, S., AND AL. Practical black-box attacks against deep learning systems using adversarial examples. *arXiv preprint arXiv:1602.02697* (2016).

[22] PAPERNOT, N., MCDANIEL, P., AND GOODFELLOW, I. J. Transferability in machine learning: from phenomena to black-box attacks using adversarial samples. *CoRR abs/1605.07277* (2016).

[23] PAPERNOT, N., MCDANIEL, P., JHA, S., FREDRIKSON, M., CELIK, Z. B., AND SWAMI, A. The Limitations of Deep Learning in Adversarial Settings. In *Proceedings of the 1st IEEE European Symposium in Security and Privacy (EuroS&P)* (2016).

[24] PAPERNOT, N., MCDANIEL, P., SINHA, A., AND WELLMAN, M. Towards the science of security and privacy in machine learning. *arXiv preprint arXiv:1611.03814* (2016).

[25] PAPERNOT, N., MCDANIEL, P., WU, X., JHA, S., AND SWAMI, A. Distillation as a defense to adversarial perturbations against deep neural networks. *Proceedings of the 37th IEEE Symposium on Security and Privacy (S&P)* (2015).

[26] ROSENBAUM, P. R. An exact distribution-free test comparing two multivariate distributions based on adjacency. *Journal of the Royal Statistical Society B*, 2005.

[27] ROZSA, A., GUNTHER, M., AND BOULT, T. E. Towards Robust Deep Neural Networks with BANG. *ArXiv e-prints* (Nov. 2016).

[28] SHEN, S., TOPLE, S., AND SAXENA, P. Auror: defending against poisoning attacks in collaborative deep learning systems. In *Proceedings of the 32nd Annual Conference on Computer Security Applications, ACSAC 2016, Los Angeles, CA, USA, December 5-9, 2016* (2016), pp. 508–519.

[29] SRNDIC, N., AND LASKOV, P. Practical evasion of a learning-based classifier: A case study. In *2014 IEEE Symposium on Security and Privacy, SP 2014, Berkeley, CA, USA, May 18-21, 2014* (2014), pp. 197–211.

[30] SZEGEDY, C., ZAREMBA, W., SUTSKEVER, I., BRUNA, J., ERHAN, D., GOODFELLOW, I., AND FERGUS, R. Intriguing properties of neural networks. In *Proceedings of the 2014 International Conference on Learning Representations* (2014), Computational and Biological Learning Society.

[31] SZÉKELY, G. J., AND RIZZO, M. L. Energy statistics: A class of statistics based on distances. *Journal of Statistical Planning and Inference 143*, 8 (2013), 1249 – 1272.

[32] WANG, B., GAO, J., AND QI, Y. A theoretical framework for robustness of (deep) classifiers under adversarial noise. *CoRR abs/1612.00334* (2016).

[33] XU, W., QI, Y., AND EVANS, D. Automatically evading classifiers. In *Proceedings of the 2016 Network and Distributed Systems Symposium* (2016).